

Gimp Tutorials

From the Source...

12/23/2017

Hisham Ismail

Chapter 1: Quickies

Intention

So you installed GIMP on your computer, congratulations! GIMP is a very powerful image manipulation software, but don't let that intimidate you. Even if you don't have time to learn advanced computer graphics, GIMP can still be a very useful and handy tool for quick image modifications.

It is my hope that these few examples will help to solve those small, quick modifications that you may need to apply to an image. Hopefully this will lead to learning even more powerful image editing capabilities that GIMP is capable of as well.

For quick access, these are the four main points I'll cover in this quick tutorial:

- [Changing the Size \(Dimensions\) of an Image \(Scale\)](#)
- [Changing the Size \(File size\) of a JPEG](#)
- [Crop an Image](#)
- [Rotate or Flip an Image](#)

In keeping with the spirit of the predecessor to this page, I will be using images from the Astronomy Picture of the Day (APOD), provided by NASA.

All you need to know to follow these quick examples is to be able to find your image and open it.

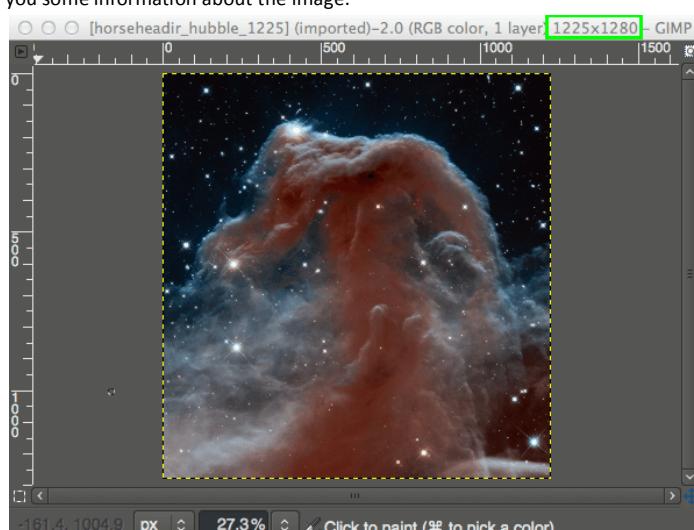
Changing the Size (Dimensions) of an Image (Scale)

It's a common problem that you may have an image that is too large for a particular purpose (embedding in a webpage, posting somewhere online, or including in an email for instance). In this case you will often want to *scale* the image down to a smaller size more suitable for your use.

This is a very simple task to accomplish in GIMP easily.

The image we'll be using to illustrate this with is The Horsehead Nebula in Infrared.

When you first open your image in GIMP, chances are that the image will be zoomed so that the entire image fits in your canvas. The thing to notice for this example is that by default the window decoration at the top of GIMP will show you some information about the image.

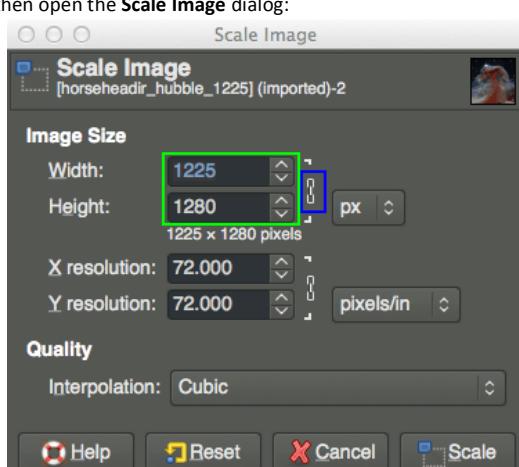


Notice that the information at the top of the window shows the current pixel dimensions of the image (in this case, the pixel size is 1225x1280).

To resize the image to new dimensions, we need only invoke the **Scale Image** dialog:

Image → Scale Image...

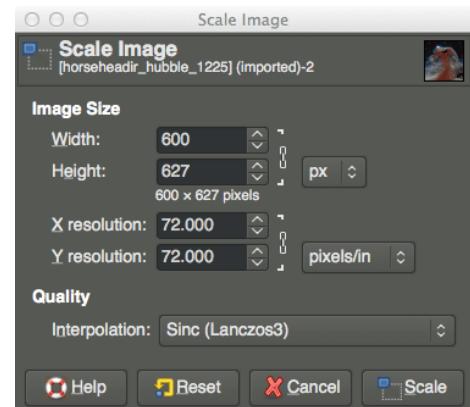
This will then open the **Scale Image** dialog:



In the **Scale Image** dialog, you'll find a place to enter new values for **Width** and **Height**. If you know one of the new dimensions you'd like for the image, fill in the appropriate one here.

You'll also notice a small chain just to the right of the **Width** and **Height** entry boxes. This icon shows that the **Width** and **Height** values are locked with respect to each other, meaning that changing one value will cause the other to change in order to keep the same aspect ratio (no strange compression or stretching in the image).

For example, if you knew that you wanted your image to have a new width of 600px, you can enter that value in the **Width** input, and the **Height** will automatically change to maintain the aspect ratio of the image:



As you can see, entering **600px** for the width automatically changes the height to **627px**.

Also notice I have shown a different option under **Quality** → **Interpolation**. The default value for this is *Cubic*, but to retain the best quality it would better to use *Sinc (Lanczos3)*.

If you want to specify a new size using a different type of value (other than Pixel size), you can change the type by clicking on the "px" spinner:



A common use for this could be if you wanted to specify a new size as a percentage of the old one. In this case you could change to "percent", and then enter 50 in either field to scale the image in half.

Once you are done scaling the image, don't forget to export the changes you've made:

File → Export...

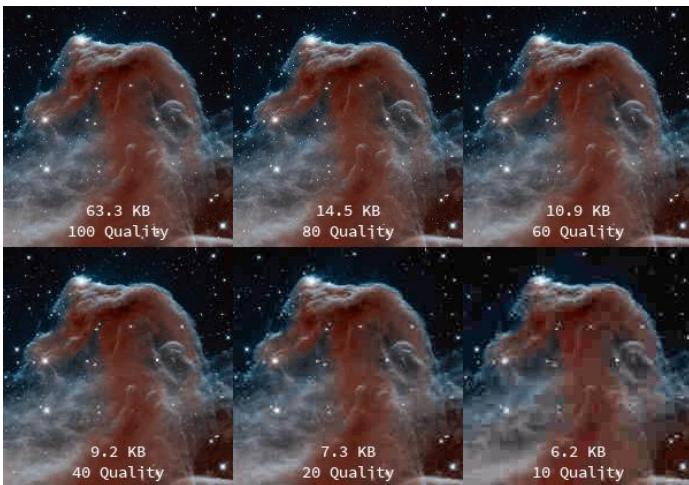
to export as a new filename, or:

File → Overwrite {FILENAME}

Changing the Size (File size) of a JPEG

You can also modify the file size of an image when exporting it to a format like JPEG. JPEG is a *lossy* compression algorithm, meaning that when saving images to the JPEG format, you will sacrifice some image quality to gain a smaller file size.

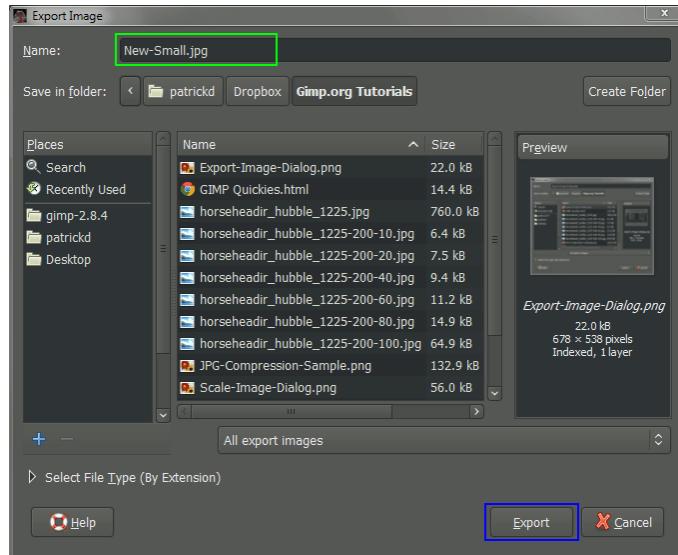
Using the same Horsehead Nebula image from above, I have resized it to 200px wide (see above), and exported it using different levels of JPEG compression:



As you can see, even at a quality setting of 80, the image is significantly smaller in file size (77% size reduction), while the image quality is still quite reasonable. When you've finished any image modifications you are doing, and are ready to export, simply invoke the export dialog with:

File → Export...

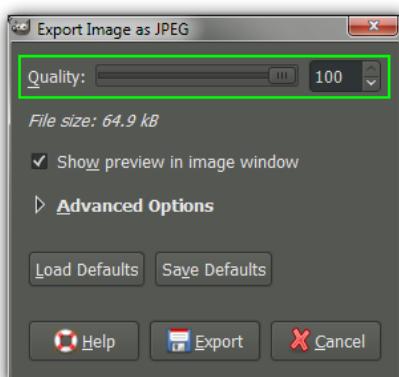
This will invoke the **Export Image** dialog:



You can now enter a new name for your file here. If you include the filetype extension (in this case, .jpg), GIMP will automatically try to export in that file format for you. Here I am exporting the image as a JPEG file.

You can also navigate to a new location on your computer through the **Places** pane, if you need to export the file to a different location. When you are ready to export the image, just hit the **Export** button.

This will then bring up the **Export Image as JPEG** dialog, where you can change the quality of the export:



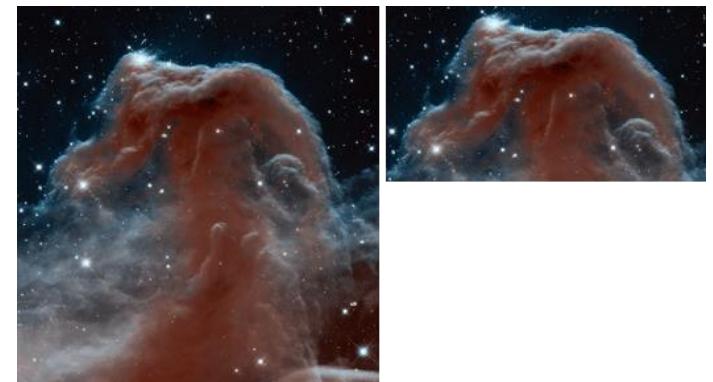
From this dialog you can now change the quality of the export. If you also have the "Show preview in image window" option checked, the image on the canvas will update to reflect the quality value you input. This will also enable the "File size:" information to tell you what the resulting file size will be. (You may need to move some windows around to view the preview on the canvas in the background).

When you are happy with the results, hit the **Export** button to export.

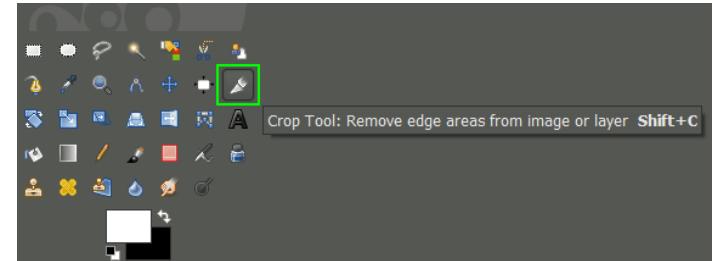
Crop an Image

There are numerous reasons you may want to crop an image. You may want to remove useless borders or information for aesthetic reasons, or you may want the focus of the final image to be of some particular detail for instance.

In a nutshell, cropping is just an operation to trim the image down to a smaller region than what you started with:



The procedure to crop an image is straightforward. You can either get to the **Crop Tool** through the tools palette:



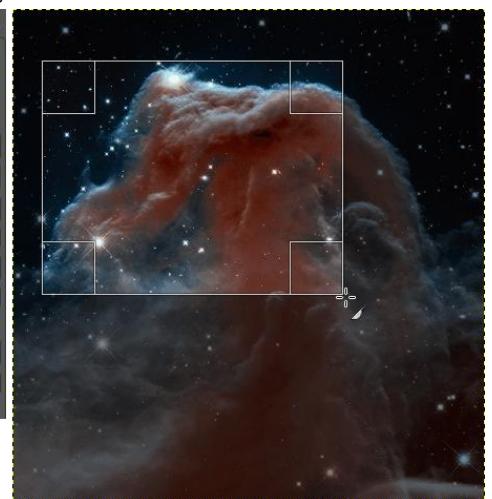
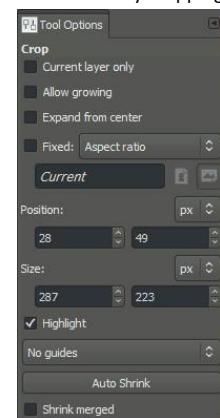
Crop Tool on the Tools Palette.

Or you can access the crop tool through the menus:

Tools → Transform Tools → Crop

Once the tool is activated, you'll notice that your mouse cursor on the canvas will change to indicate the **Crop Tool** is being used.

Now you can Left-Click anywhere on your image canvas, and drag the mouse to a new location to highlight an initial selection to crop. You don't have to worry about being exact at this point, as you will be able to modify the final selection before actually cropping.

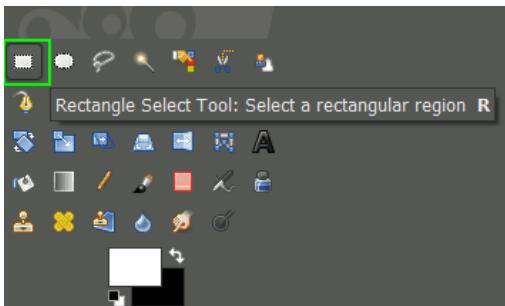


After making the initial selection of a region to crop, you'll find the selection still active. At this point hovering your mouse cursor over any of the four corners or sides of the selection will change the mouse cursor, and highlight that region.

This allows you to now fine-tune the selection for cropping. You can click and drag any side or corner to move that portion of the selection.

Once you are happy with the region to crop, you can just press the **"Enter"** key on your keyboard to commit the crop. If at any time you'd like to start over or decide not to crop at all, you can press the **"Esc"** key on your keyboard to back out of the operation.

Another way to crop an image is to make a selection first, using the **Rectangle Select Tool**:



Or through the menus:

Tools → Selection Tools → Rectangle Select

You can then highlight a selection the same way as the **Crop Tool**, and adjust the selection as well. Once you have a selection you like, you can crop the image to fit that selection through:

Image → Crop to Selection

Rotate and/or Flip an Image

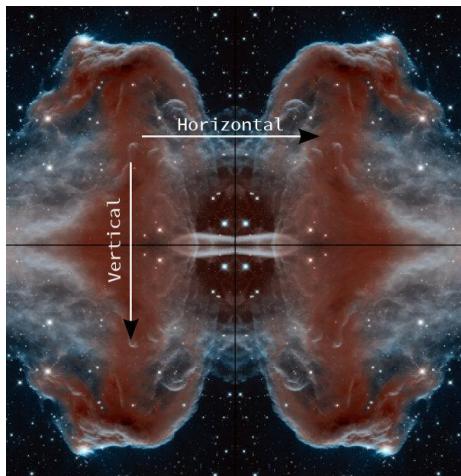
There may be a time that you would need to rotate an image. For instance, you may have taken the image with your camera in a vertical orientation, and for some reason it wasn't detected by GIMP as needing to be rotated (GIMP will normally figure this out for you, but not always).

There may also be a time that you'd like to flip an image as well. These commands are grouped together under the same menu item:

Image → Transform

Flip an Image

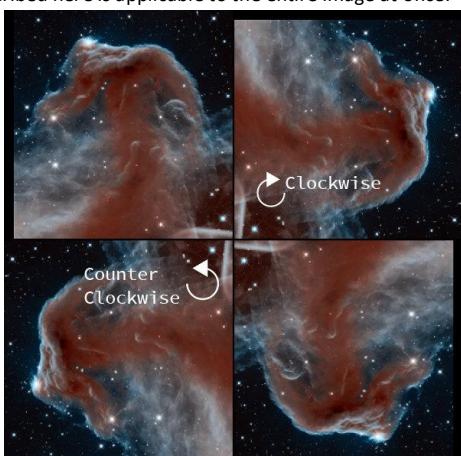
If you want to flip your image, the **Transform** menu offers two options, **Flip Horizontally**, or **Flip Vertically**. This operation will mirror your image along the specified axis. For example, here are all of the flip operations shown in a single image:



Rotate an Image

Image rotation from the **Transform** menu is constrained to either 90° clockwise/counter-clockwise, or 180°.

Don't mis-interpret this to mean that GIMP cannot do arbitrary rotations (any angle). Arbitrary rotations are handled on a per-layer basis, while the image rotation described here is applicable to the entire image at once.



Chapter 2: Simple floating Logo

Intention

This tutorial is intended to introduce you to a few simple commands, and some concepts in order to create a logo that appears to be floating above a background, like this:



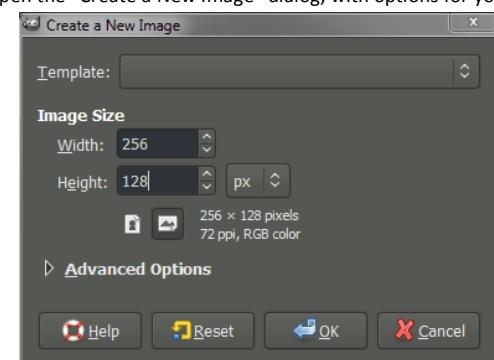
The concepts are ones that you'll likely come across multiple times while working in graphics processing. Layer masks are used to isolate a part of an image, thus allowing it to be placed over a random background for instance. The addition of a drop-shadow effect to make an object appear to be floating over the background is another example.

Getting Started

Create a new image of appropriate size for your logo:

File → New...

This will open the "Create a New Image" dialog, with options for you to specify:

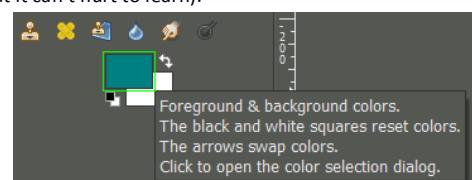


You can make this new image any dimensions you want, but for this tutorial I am going to specify a Width of 256 px, and a Height of 128 px. I haven't specified any other options. When you're ready, hit "OK" to create the new image.

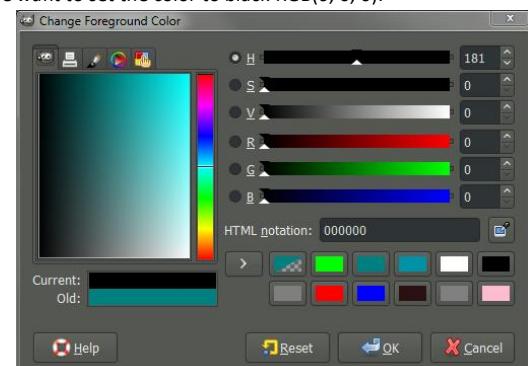
You'll be presented with the new image on your canvas. Chances are it will be a pure white image at this point (it may be a different color depending on how your GIMP is setup to handle new images - if it is, don't worry).

Fill the New Image with Black

The first thing we are going to do is fill our new image with black. The first step to doing so is to make sure that the Foreground Color is appropriately set. Click on the foreground color in the Color area to bring up the "Change Foreground Color" dialog (if your foreground color is already black you don't have to do this step, but it can't hurt to learn):

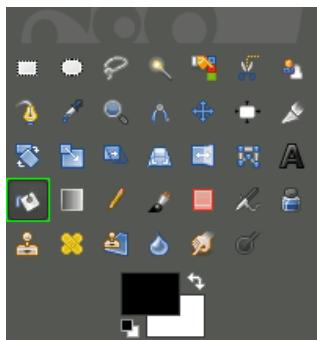


The "Change Foreground Color" dialog allows you to now set the foreground color. We want to set the color to black RGB(0, 0, 0):



With the foreground color set, we can now use the Bucket Fill Tool to fill in our image:

Tools → Paint Tools → Bucket Fill

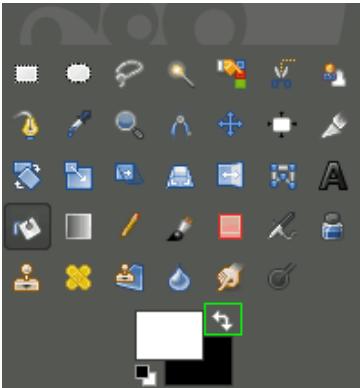


Once the tool is activated, your cursor should appear as to the left. To fill the layer you need only click on the image area at this point. Your image should now fill with black.

Adding Some Text

Now we want to add text to our image to create our logo with. To see what we're doing, though, will require us to change the foreground color to something other than black (black text on black background doesn't show up so well).

Now, you can follow the above procedures again to set the foreground color. If your background color is already white, though, you can quickly swap foreground/background colors using the arrows:

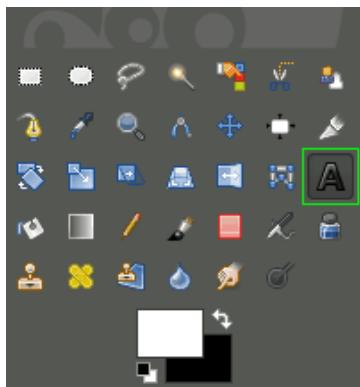


Swap Foreground/Background quickly.

You can also use the keyboard shortcut "X" to swap the colors.

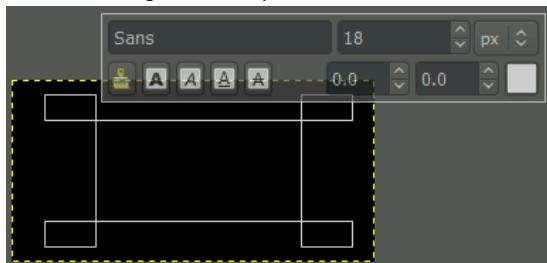
With the foreground color set to white, we can now use the **Text Tool** to add some text to our image:

Tools → Text

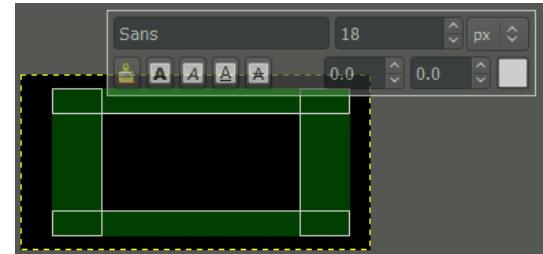


We can now draw a box on our canvas (image) to hold the text. You can click on the canvas where you'd like the top-left corner of your box to be, and drag the mouse down to the bottom right corner. You don't have to worry about being exact at this point, because you can adjust the boundaries of the box after the fact.

This is what you should see on your canvas after clicking and dragging from the top-left to the bottom-right to define your text box:



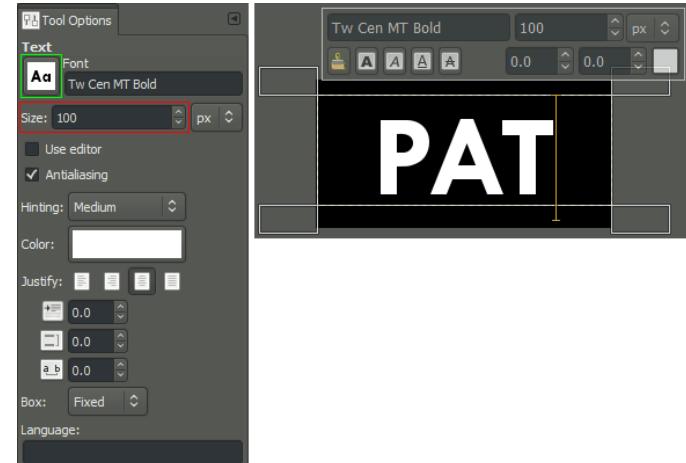
If you'd like to re-size the box for some reason, you can now click and drag in any of the green areas shown below:



Your text will go into the black box inside the green areas shown above.

Once the text boundary box is sized appropriately, we can just type some text.

In my case, I'll use my name:

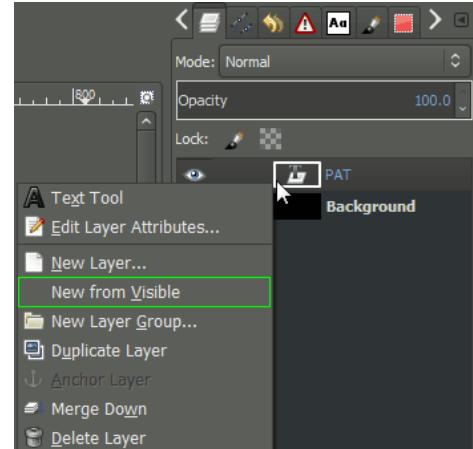


Text Tool Options (left), canvas view (right).

Chances are when you first start entering text, it will be very small on your canvas. So let's have a look at some options on the **Text Tool Options** palette (left, above).

If you want to make your text appear bigger, you can change the **Size** in the field shown. For instance, here I've chosen to set my **Size** to 100px.

You may also not like the font that is chosen by default. In that case, we can change the **Font** to something better by clicking the icon. This will open a dropdown to scroll through all the fonts that GIMP knows about on your system. You can see in my example that I've changed the font to "Tw Cen MT Bold".

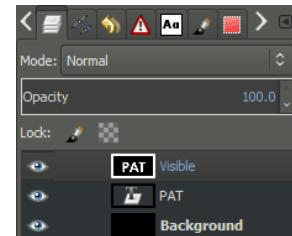


Once we've gotten the text how we want it, we can now create a new layer from all the visible layers so far (the text layer, and the black background layer). On your **Layers** tab, right click on the text layer we just made, and choose "New from Visible".

Alternatively, you can also create a new layer from visible using the menu:

Layer → New from Visible

At this point, our layer palette will have three layers on it, the background, the text ("PAT"), and our new layer "Visible":



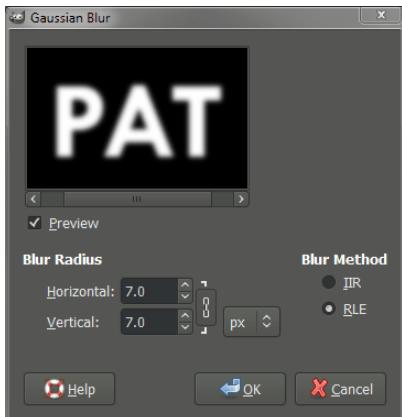
Notice that there is a white border around the “Visible” layer. This indicates that this layer is currently active, so that any operations we perform will apply to this layer.

Which is good, because we are about to blur this new layer a bit!

To apply a slight Gaussian blur to this layer, we simply invoke the command through the menu:

Filters → Blur → Gaussian Blur...

This will invoke the **Gaussian Blur** dialog, where we can specify how much blur we want to apply:



The defaults were pretty good, but I wanted just a tad more blur, so I increased the **Blur Radius** to 7. When you’re done, just hit “OK”.

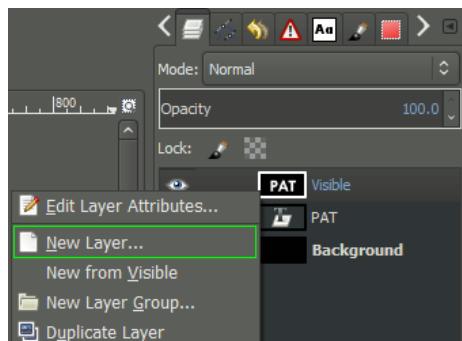
Adding Some Color

Now that we have our text done, it’s time to add a splash of color!

We are going to add a new layer to our image first:

Layer → New Layer...

Or by Right-Clicking on the “Visible” layer in the layer palette, and choosing “New Layer...” from the context menu:



The “Create a New Layer” dialog will appear - it doesn’t matter what it gets filled with, so you can leave it at whatever **Layer Fill Type** it’s set at (White by default I believe). Hit **OK** to create the new layer.

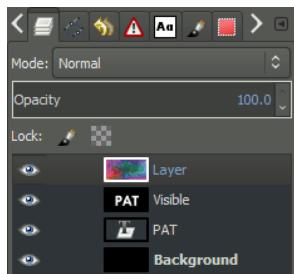
We are now going to fill this new layer with some color to add some interest.

To do this we will use the **Plasma** plugin:

Filters → Render → Clouds → Plasma...

I just left the default values and hit **OK**, but feel free to fiddle with the values.

Our layers now look like this:



Here is what my canvas looks like right now (with the plasma layer on top and visible):

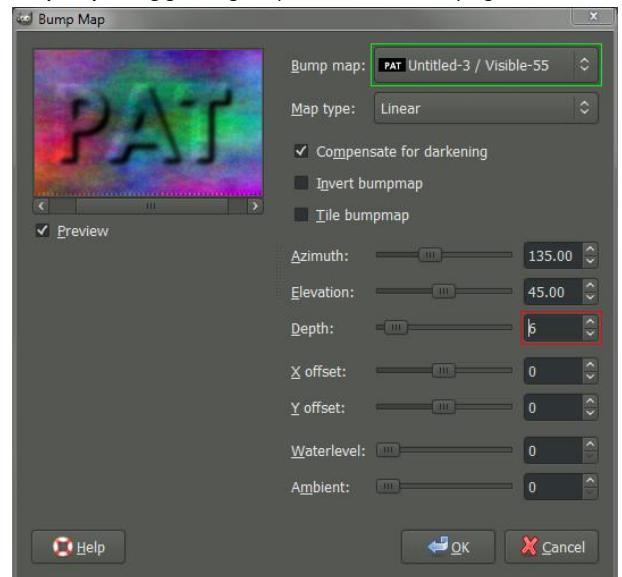


Bump Mapping

Now we’re going to use the text we created earlier to generate a fake 3D shape on this plasma layer. The process is known as “bump mapping”. I won’t get into the technical details of how this works, as it is best seen rather than explained. Open the **Bump Map** dialog through the menu:

Filters → Map → Bump Map...

The **Bump Map** dialog gives a good preview of what the plugin does:



To get it working correctly, this plugin requires that you properly point to the source for the bump mapping. In our case the source is the text layer we created earlier (the layer was named “Visible”). So we’ll click on the spinner for the **Bump map**, and choose our “Visible” layer from the list.

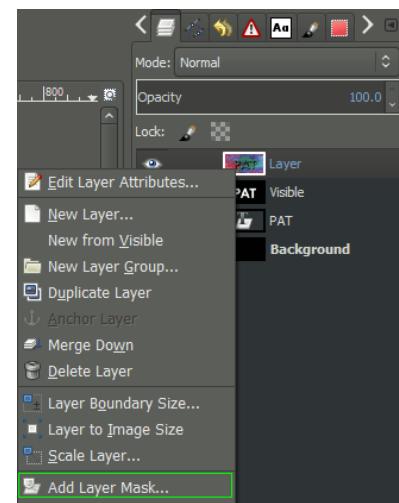
As before, feel free to play with the options. The only one that I changed was the **Depth** to increase the illusion of depth (I finally set the value to 6 in my example). Once it looks good, we’ll hit the **OK** button to apply it to the layer.

Apply a Layer Mask

Now we are going to use a **Layer Mask** to isolate our bump mapped text. First we need to add a **Layer Mask** to the plasma layer:

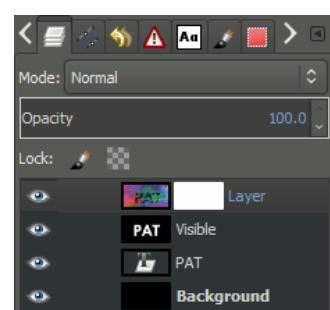
Layer → Mask → Add Layer Mask...

Or Right-Click on the plasma layer and choose “Add Layer Mask...” from the context menu:



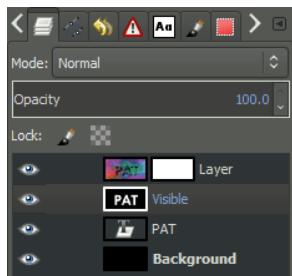
When the “Add a Mask to the Layer” dialog comes up, set the **Initialize Layer Mask to: White** (full opacity).

Once you’ve added a mask to the plasma layer, your layers should now look like this:



Remember, you can tell which layer (or mask) is active by noticing which one has the white border around it. The layers above show that the plasma layers mask is active (there is a white border around the white mask, so it's hard to notice, but *no other* layer/mask has a white border).

We are going to copy the "Visible" layer, and paste it into the layer mask for the plasma layer. So first, Left-Click on the "Visible" layer in the layers palette to activate it:

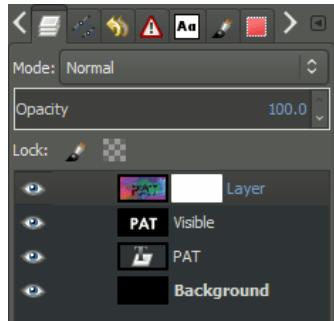


Remember, the white border will indicate the layer is active.

With the layer active, we want to now copy it:

Edit → Copy

Then we want to make the plasma layer mask active by Left-Clicking on the mask:

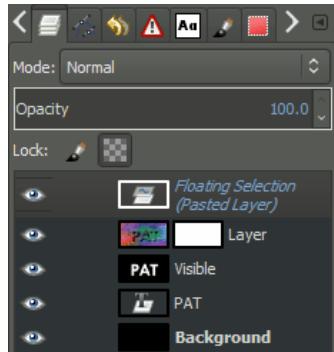


Plasma layer mask now active again.

With the mask active again, we now want to paste the "Visible" layer back into the image:

Edit → Paste

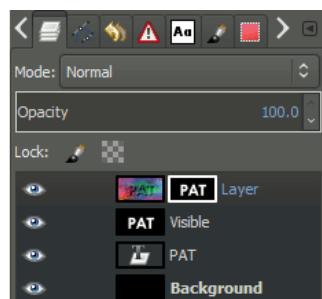
This will now insert a *Floating Selection (Pasted Layer)* into your image:



To get this *Floating Selection* into the mask, we need to **Anchor** it:

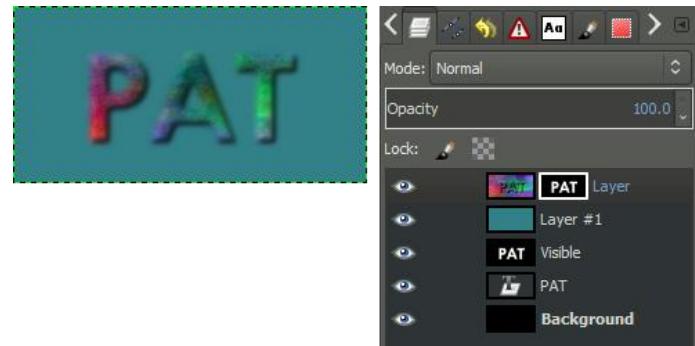
Layer → Anchor Layer

This will **Anchor** the selection down onto the mask. Our image and layers should now look something like this:



We may now want to add a different colored background to help us fine-tune the results we have so far. Add a new layer to the image as we did when [creating the plasma layer](#), and place it below the plasma layer. (You can click and drag layers to change their order in the palette).

Pick an interesting background color and fill the new layer with this color. The layers should now look like this:



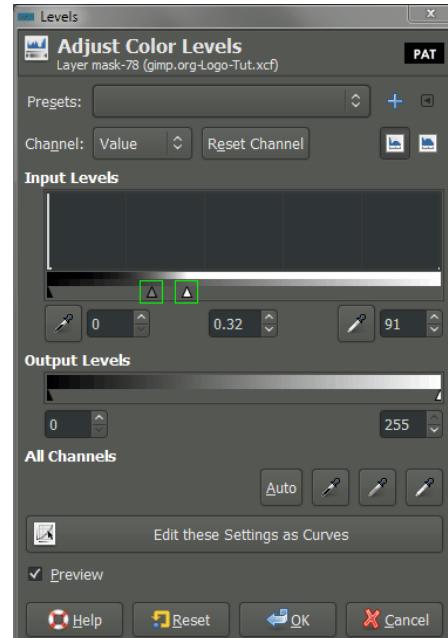
It doesn't look bad, but we can perhaps tighten up the results by adjusting the mask a bit to clarify the edges of the text.

Adjusting the Levels

I want to clean up the edges of the text with what we have so far. Right now, the mask being used on the plasma layer is a copy of the gaussian blurred text. To make it sharper, we are going to adjust the levels on the mask for that layer. To do this, we first make sure the layer mask is active by clicking on it. Then we can open the **Adjust Color Levels** dialog through the menu:

Colors → Levels...

With the **Adjust Color Levels** dialog, we now want to sharpen up the edges of the mask a little bit:



What we want to do is adjust the **Gamma** and **White point** sliders. I started by dragging the **White point** slider down to increase the prominence of the plasma layer, then pushed the **Gamma** up to emphasize it more. (If you're following along, you can also just copy my values from above).

The trick is to increase the definition of the edges of the text, without going too far and causing it to look very jagged (aliased). Play with the settings to see how they affect your results. Here is what my plasma layer looks like after applying the above levels to the mask:



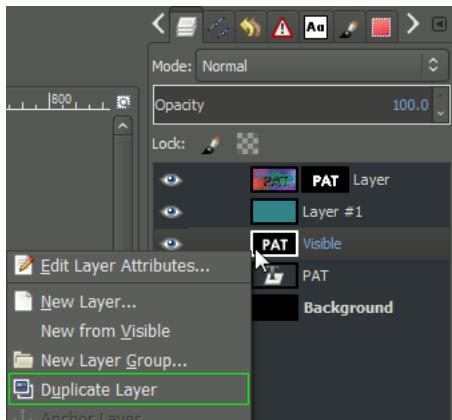
Creating a Drop Shadow

Now we may want to get a little fancier and add an effect of a drop shadow behind the logo to make it seem as if it's floating above the background. We've already created what we need to generate this effect, we just need to move a couple of things around to do so.

Make a copy of your "Visible" layer that had your original blurred text on it. Select the layer first to activate it, then you can do:

Layer → Duplicate Layer

Or Right-Click on the "Visible" layer, and choose "Duplicate Layer":



This will create a new layer called "Visible copy". Move this layer above your background color layer to just beneath your plasma layer as shown (you can Left-Click and drag the layer in the palette):



Click and drag the "Visible copy" layer to beneath the plasma layer. This layer will become our shadow, but we need to modify a couple of things first. We'll first invert the colors of the layer to make the text black using:

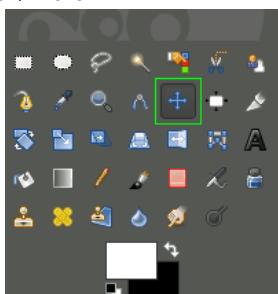
Colors → Invert

Then we need to change the layer so that all of the white areas will be transparent. This can be found in the menu:

Layer → Transparency → Color to Alpha...

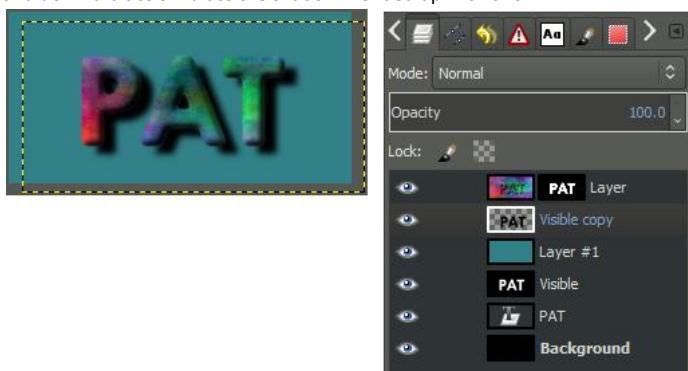
The layer should now have black text over a transparent background. We'll now just want to shift this layer a bit to simulate a height by offsetting it down and to the right a bit. To do this we can use the **Move Tool**:

Tools → Transform Tools → Move



Activate the **Move Tool**.

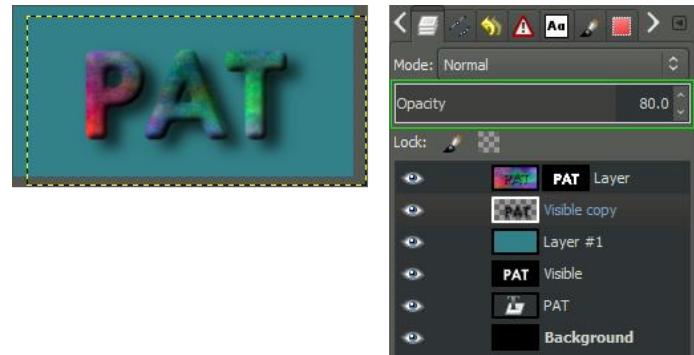
If we hold down **Shift** and click on the canvas, we restrict the **Move Tool** to modifying only the active layer (our shadow layer). Drag the layer to the right and down a bit to simulate the shadow. I ended up with this:



Shadow layer shifted to the right and down a bit.

At this point we can make it a bit more fancy by adding a **Gaussian Blur** to the shadow to spread it out a little more. We can also modify the layer **Opacity**, adjusting it to let the background show through a bit as well.

Here is the final state of my image, where I applied a **Gaussian Blur** with a 10px radius, and adjusted the shadow layer **Opacity** down to 80%:



The End

Here is my final floating logo image when everything is done:



The neat thing about our process is that we can now use any background we want behind the image, and the effects and shadow will still be there:



Or we could save it as a PNG file with no background at all, thus allowing whatever background there is to show through:



So that's it! Go on and have fun making floating logos.

Chapter 3: Circle Images

Making a circle-shaped image

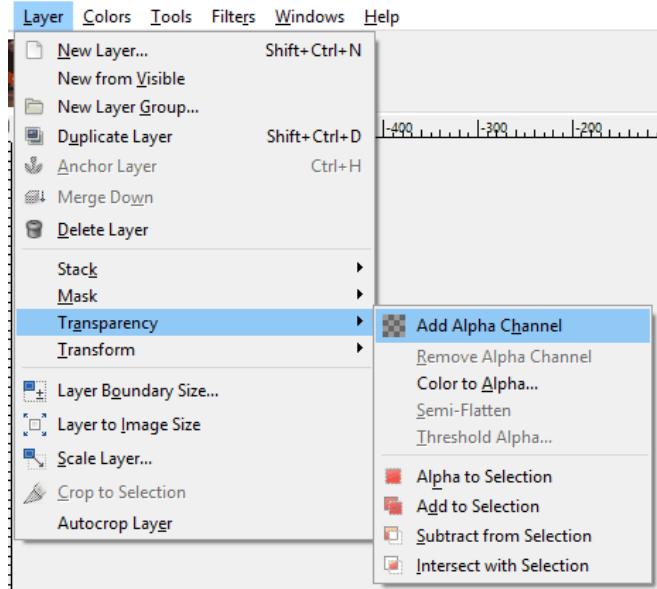
There are no circular images. There are only rectangular images. But there can be images where corners are transparent, so that only a circle shows.

In light of this:

- First, make sure that your layer has an “alpha channel”.

Layer → Transparency → Add alpha channel

if it's grayed out it means you already have one



- Create a circular selection with the “Ellipse select tool” (the 2nd one in the toolbox).
- Use the “Tool options” dialog

Windows → Dockable dialogs → Tool options



- If you want a true circle, use the *Fixed* option: select *Aspect ratio* and enter 1:1.
- Depending on what kind of marks you have, you can use:
 - The diagonal framing (default): click on one corner, drag across a full diagonal and release at the opposite corner,
 - The radial framing (check *Expand from center* in the Tool options): click on the center, drag across a half diagonal release on a corner.
- If the selection isn't perfect on the first try, you can move it (click around the middle) or extend it (click inside, near a border or a corner).
- Once you have the required selection, invert the selection (*Select → Invert*, or *Ctrl-I*) so that everything is selected, except your circle.
- Erase the selection (*Edit → Clear* or *[Delete]* key). You should have your central circle left, surrounded by a checkerboard pattern. (this checkerboard is not part of the image, it just indicates the transparent parts of the image).
- You can reduce the checkerboard to the minimum by auto-cropping the image (*Image → Autocrop image*)
- Last, save the image in a format that supports transparency, like PNG (JPEG doesn't support transparent images...)
- If you are going to work further on the picture, save it as XCF (Gimp native format).

Chapter 4: Layer Masks

Intention

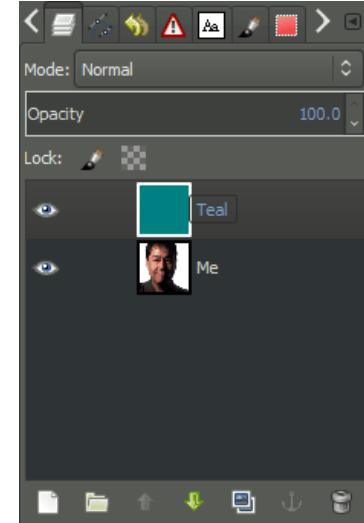
Layer masks are a fundamental tool in image manipulations. They allow you to selectively modify the opacity (transparency) of the layer they belong to. This differs from the use of the layer **Opacity** slider as a mask has the ability to *selectively* modify the opacity of different areas across a single layer.

This modification of a layer's transparency through a mask is non-destructive to the layer itself.

This flexibility to define the opacity of different areas of a layer is the basis for more interesting image manipulation techniques such as selective coloring and luminosity masking.

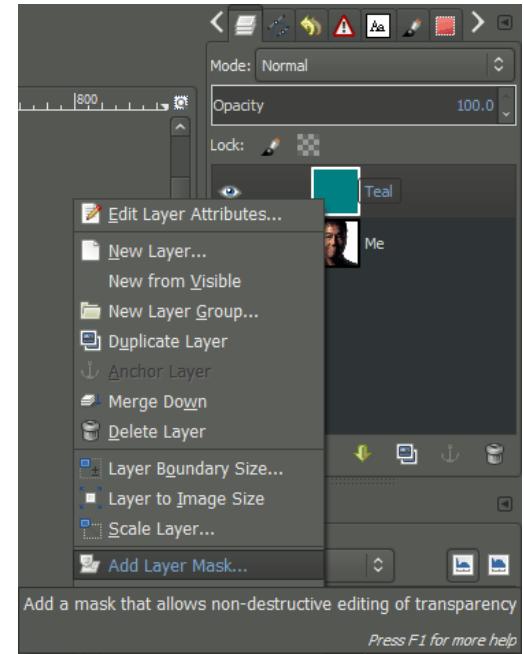
Adding a Mask to a Layer

Layer masks need to be added to a layer before they can be used. The process for adding them is simple.



For this example I will use a simple image with only two layers, as shown above. There is a base image at the bottom of the stack, and a single layer of teal over it. The teal layer is the active layer (look for the white border), and the one which we will add a layer mask to.

Right-Click on the layer you want to add a mask to (the “Teal” layer in my example), and the Context menu will show an option to **Add Layer Mask...**:

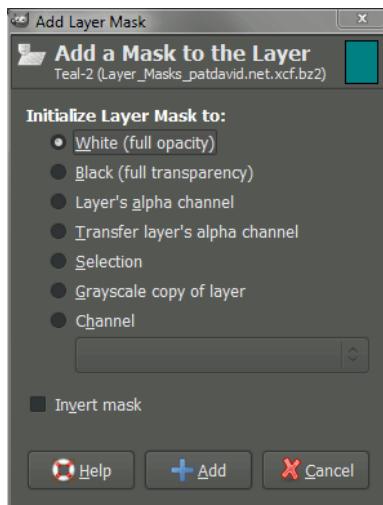


Add Layer Mask in the context menu.

You can also add a layer mask through the menus:

Layer → Mask → Add Layer Mask...

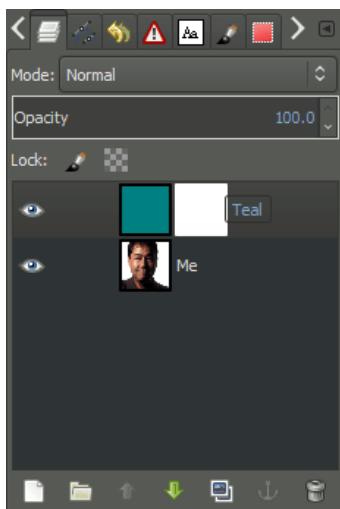
This will then bring up the “Add a Mask to the Layer” dialog with some options:



Add mask options dialog.

There are many options for initializing the Layer Mask. Notice that the first option is to set the entire mask to White, which will result in full opacity on the layer (no transparency from the mask). The option to initialize to Black shows that the mask will make the entire layer fully transparent.

For the purposes of this tutorial, we will let the mask initialize to **White** (full opacity). You should notice a change in your Layers dialog now that shows the layer mask thumbnail to the right of the layer it applies to (in this case the "Teal" layer):



Layers dialog with mask applied to Teal layer.

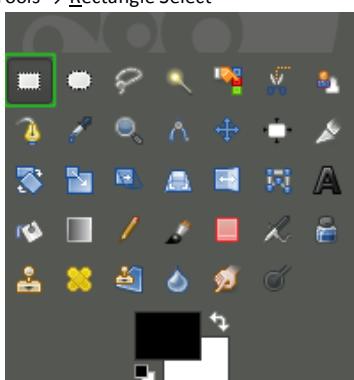
The layer mask has now been added to the "Teal" layer. It is also active (there is a white border around the thumbnail in the dialog, but is not visible due to the mask being white as well) and ready for modification.

Modifying a Layer's Transparency with the Mask

At this point any operations performed on the canvas will apply to the mask and not to any layers themselves. To illustrate how masks can affect its layers transparency, let's paint!

I am going to use the **Rectangle Select** tool to select roughly the top third of the image, and I'll fill this selection with black.

Tools → Selection Tools → Rectangle Select



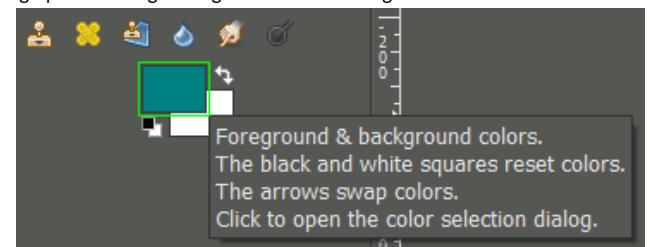
Activating the **Rectangle Select** tool

Using the **Rectangle Select** tool, I'll select roughly the top third of the image:



Top third of the image selected.

I want to fill this selection with black, but before I do I need to make sure that my foreground color is black. Click on the foreground color in the **Color area** to bring up the "Change Foreground Color" dialog:



Click the foreground color to change.

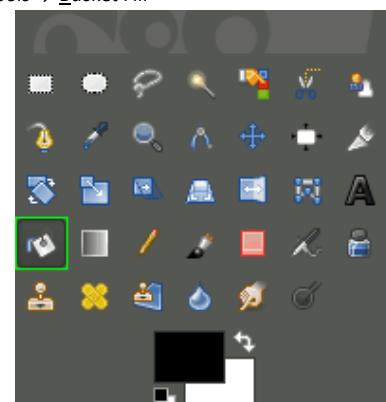
The "Change Foreground Color" dialog allows you to set the foreground color. For this example set the color to black, RGB (0, 0, 0):



Change the color to black.

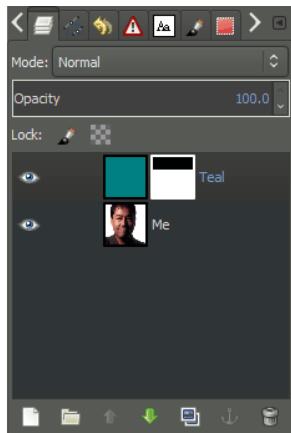
With the foreground color set, you can now use the **Bucket Fill Tool** to fill in the selection.

Tools → Paint Tools → Bucket Fill

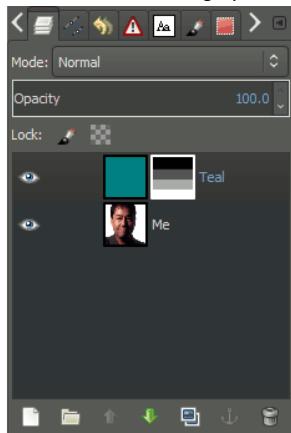
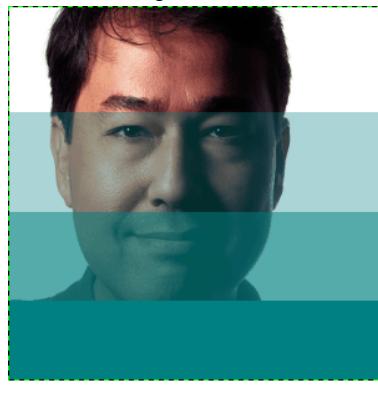


Activating the **Bucket Fill** tool

You can now click inside the selection to fill it with the foreground color (black). As soon as you do, you'll be presented with a new view of your image on the canvas:



As you can see, filling the selected portion of the layer mask with black resulted in that area having 100% transparency, showing the layer below it. If you **Rectangle Select** a different area of the mask, you can fill it in with a different shade of gray to produce a variable opacity. For example, I will select a few different regions of the mask, and fill it with different levels of gray:



If you examine the layer mask, you'll see that there are different levels of gray being applied (black to white, from top to bottom), and their value is what determines the opacity of the layer.

Selective Colorization Example

A good example of the application of layer masks is doing selective colorization of an image (selectively allowing color to show through a mostly black and white image). I'll walk through how to easily do this with an image from Mardi Gras 2013:



Mardi Gras 2013, Mobile, AL

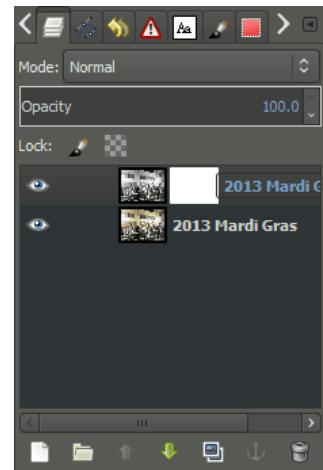
Start the process by duplicating the base image (Shift+Ctrl+D, or Right-Click layer → Duplicate Layer). From the menu:

Layer → Duplicate Layer

Then desaturate the upper layer using:

Colors → Desaturate

Following the steps above, add a layer mask to the desaturated layer and initialize it to White (full opacity). At this point, the Layers dialog should look like this:



As before, set your foreground color to black. This time, rather than filling selections, we are going to use the **Paintbrush Tool** to paint areas of the image we want the color to show through from the layer below.

I decided to paint the boy on the fence. Using the **Paintbrush Tool** I painted over his shirt and head. This allowed those colors to show through from the layer below. Here are the results after painting:



Simple Selective Colorization Example

To illustrate what was done, here is the layer mask I painted to achieve the above result:



Layer mask isolated to illustrate

Of course, I could have chosen a different color than black to create the mask. If I wanted a slightly more muted color I could have painted with a more middle gray vs. black:



Simple Selective Colorization Example (painted with gray vs. black).

Chapter 5: Basic Color Curves

Color has this amazing ability to evoke emotional responses from us. From the warm glow of a sunny summer afternoon to a cool refreshing early evening in fall. We associate colors with certain moods, places, feelings, and memories (consciously or not).

Volumes have been written on color and I am in no ways even remotely qualified to speak on it. So I won't.

Instead, we are going to take a look at the use of the **Curves** tool in [GIMP](#). Even though GIMP is used to demonstrate these ideas, the principles are generic to just about any RGB curve adjustments.

Your Pixels and You

First there's something you need to consider if you haven't before, and that's what goes into representing a colored pixel on your screen.



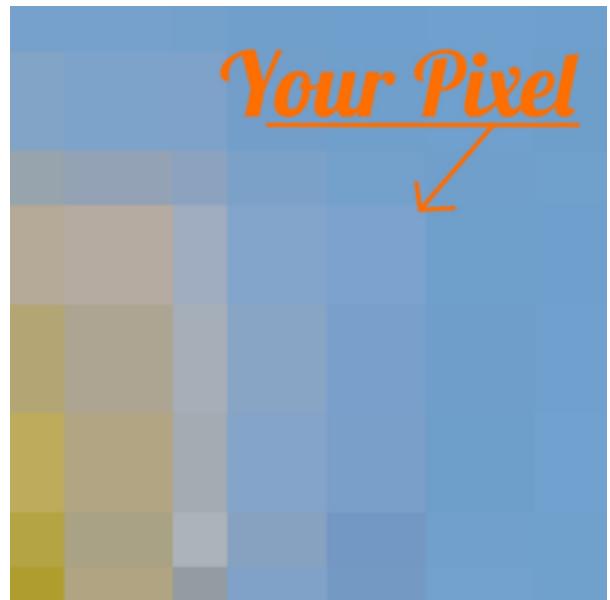
Open up an image in GIMP.



Now zoom in.



Nope - don't be shy now, zoom in more!



And there's your pixel.

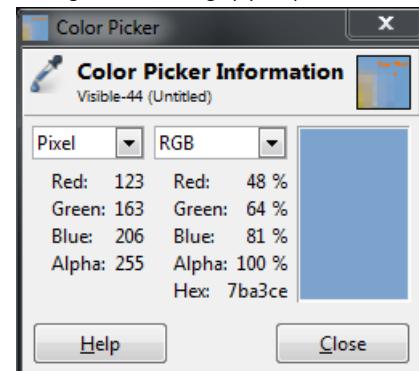
So let's investigate what goes into making your pixel.

Remember, each pixel is represented by a combination of 3 colors: **Red**, **Green**, and **Blue**. In GIMP (currently at 8-bit), that means that each RGBcolor can have a value from **0 - 255**, and combining these three colors with varying levels in each channel will result in all the colors you can see in your image.

If all three channels have a value of 255 - then the resulting color will be pure white. If all three channels have a value of 0 - then the resulting color will be pure black.

If all three channels have the same value, then you will get a shade of gray (128,128,128 would be a middle gray color for instance).

So now let's see what goes into making up your pixel:



The RGB components that mix into your final blue pixel

As you can see, there is more blue than anything else (it is a blue-ish pixel after all), followed by green, then a dash of red. If we were to change the values of each channel, but kept ratio the same between Red, Green, and Blue, then we would keep the same color and just lighten or darken the pixel by some amount.

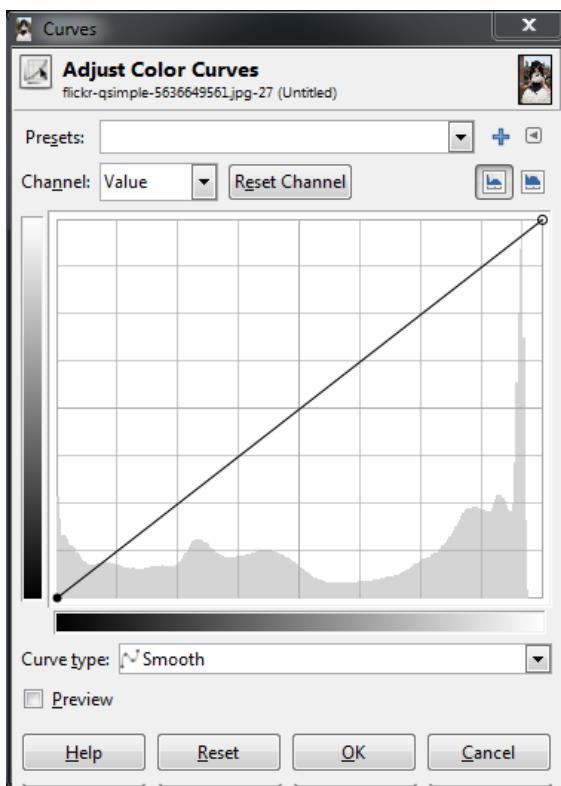
Curves: Value

So let's leave your pixel alone for the time being, and actually have a look at the **Curves** dialog. I'll be using this wonderful image by [Eric](#) from [Flickr](#).

Intentionally Removed the Image from here. It is not so important.

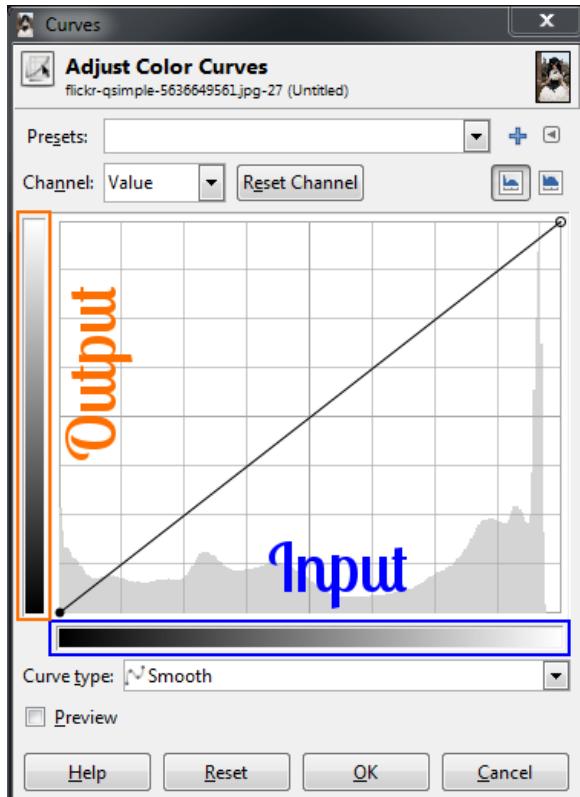
[Hollow Moon](#) by [gsimple/Eric](#) on Flickr. (cc-by-nc-sa)

Opening up my **Curves** dialog shows me the following:
Colors → Curves...



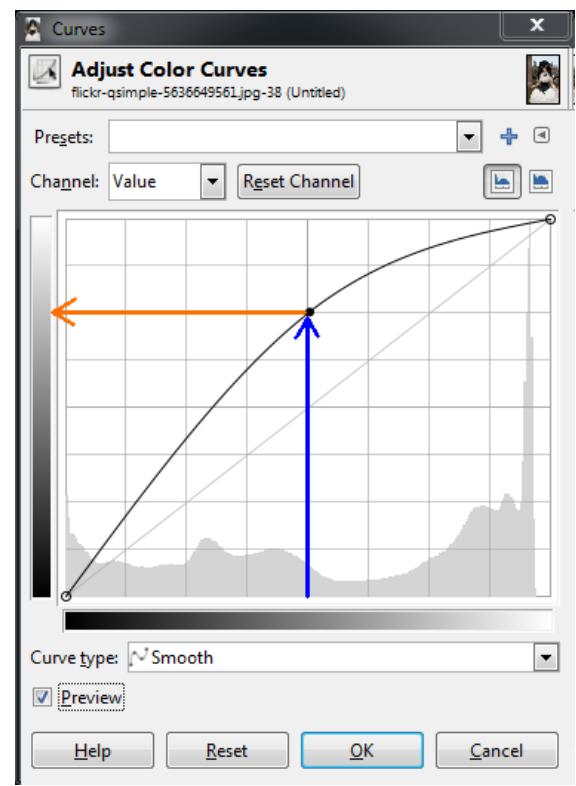
We can see that I start off with the curve for the **Value** of the pixels. I could also use the drop down for “**Channel**” to change to red, green or blue curves if I wanted to. For now let’s look at **Value**, though.

In the main area of the dialog I am presented with a linear curve, behind which I will see a histogram of the value data for the entire image (showing the amount of each value across my image). Notice a spike in the high values on the right, and a small gap at the brightest values.



What we can do right now is to adjust the values of each pixel in the image using this curve. The best way to visualize it is to remember that the bottom range from black to white represents the **current** value of the pixels, and the left range is the value to be mapped to.

So to show an example of how this curve will affect your image, suppose I wanted to remap all the values in the image that were in the midtones, and to make them all lighter. I can do this by clicking on the curve near the midtones, and dragging the curve higher in the Y direction:



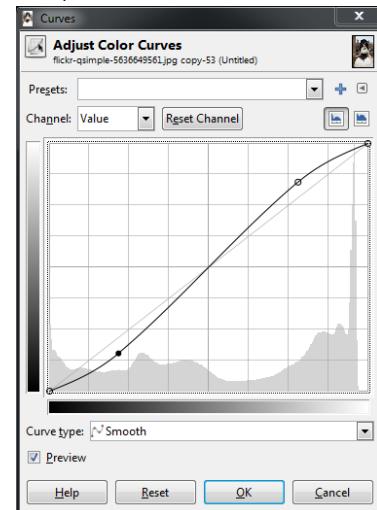
What this curve does is takes the values around the midtones, and pushes their values to be much lighter than they were. In this case, values around 128 were re-mapped to now be closer to 192.

Because the curve is set **Smooth**, there will be a gradual transition for all the tones surrounding my point to be pulled in the same direction (this makes for a smoother fall-off as opposed to an abrupt change at one value). Because there is only a single point in the curve right now, this means that all values will be pulled higher.

Intentionally Removed the Image from here. It is not so important.

The results of pushing the midtones of the value curve higher (click to compare to original).

Care should be taken when fiddling with these curves to not blow things out or destroy detail, of course. I only push the curves here to illustrate what they do. A very common curve adjustment you may hear about is to apply a slight “S” curve to your values. The effect of this curve would be to darken the dark tones, and to lighten the light tones - in effect increasing global contrast on your image. For instance, if I click on another point in the curves, and adjust the points to form a shape like so:



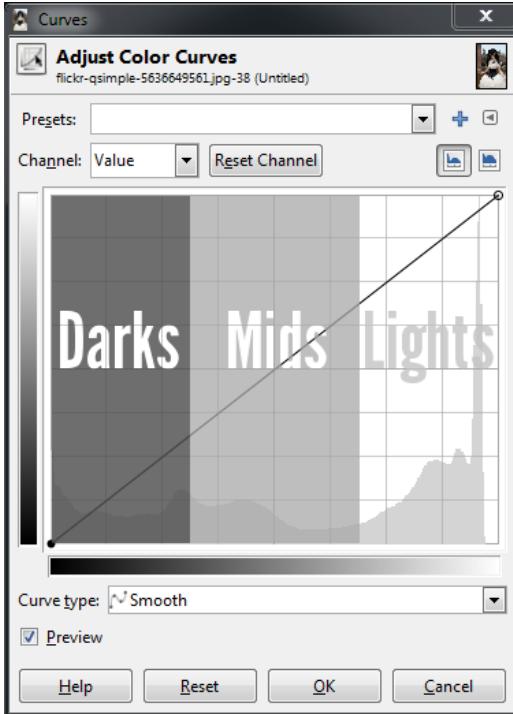
A slight “S” curve

This will now cause dark values to become even darker, while the light values get a small boost. The curve still passes through the midpoint, so middle tones will stay closer to what they were.

Intentionally Removed the Image from here. It is not so important.

Slight "S" curve increases global contrast (click for original).

In general, I find it easiest to visualize in terms of which regions in the curve will effect different tones in your image. Here is a quick way to visualize it (that is true for value as well as RGB curves):

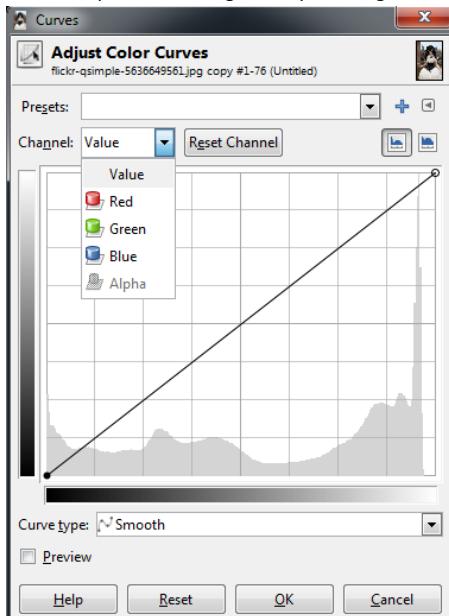


If there is one thing you take away from reading this, let it be the image above.

Curves: Colors

So how does this apply to other channels? Let's have a look.

The exact same theory applies in the RGB channels as it did with values. The relative positions of the darks, midtones, and lights are still the same in the curve dialog. The primary difference now is that you can control the contribution of color in specific tonal regions of your image.



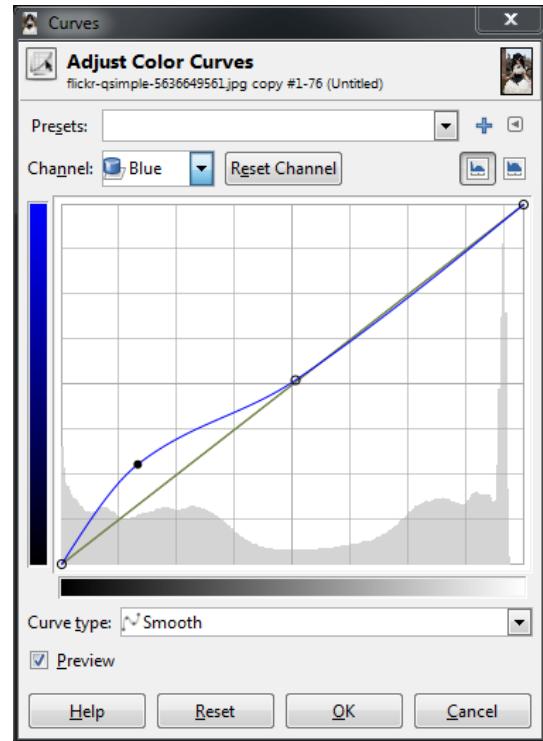
Value, Red, Green, Blue channel picker.

You choose which channel you want to adjust from the "Channel" drop-down. To begin demonstrating what happens here it helps to have an idea of generally what effect you would like to apply to your image. This is often the hardest part of adjusting the color tones if you don't have a clear idea to start with.

For example, perhaps we wanted to "cool" down the shadows of our image. "Cool" shadows are commonly seen during the day in shadows out of direct

sunlight. The light that does fall in shadows is mostly reflected light from a blue-ish sky, so the shadows will trend slightly more blue.

To try this, let's adjust the **Blue** channel to be a little more prominent in the darker tones of our image, but to get back to normal around the midtones and lighter.



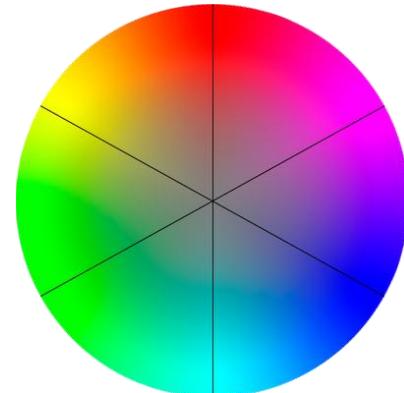
Boosting blues in darker tone

Intentionally Removed the Image from here. It is not so important.

Pushing up blues in darker tones (click for original).

Now, here's a question: If I wanted to "cool" the darker tones with more blue, what if I wanted to "warm" the lighter tones by adding a little yellow?

Well, there's no "Yellow" curve to modify, so how to approach that? Have a look at this HSV color wheel below:



The thing to look out for here is that opposite your blue tones on this wheel, you'll find yellow. In fact, for each of the Red, Green, and Blue channels, the opposite colors on the color wheel will show you what an absence of that color will do to your image. So remember:

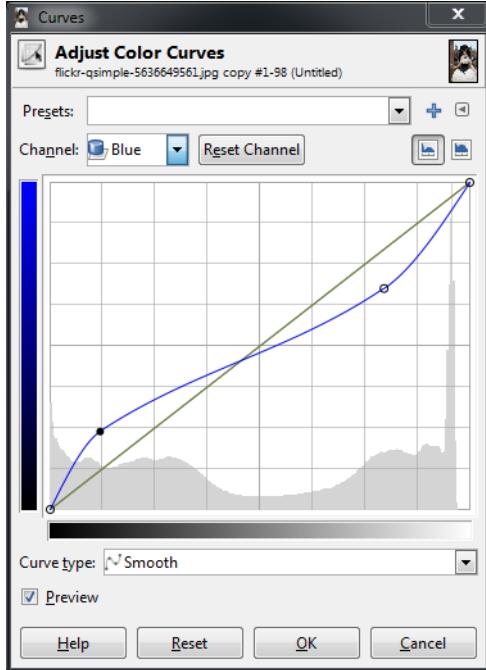
Red → Cyan

Green → Magenta

Blue → Yellow

What this means to you while manipulating curves is that if you drag a curve for blue up, you will boost the blue in that region of your image. If instead you drag the curve for blue down, you will be **removing** blues (or boosting the **Yellows** in that region of your image).

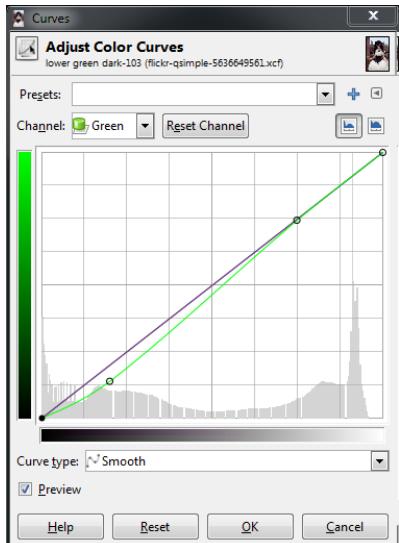
So to boost the blues in the dark tones, but increase the yellow in the lighter tones, you could create a sort of “reverse” S-curve in the blue channel:



Intentionally Removed the Image from here. It is not so important.

Boost blues in darks, boost yellow in high tones (click for original).

In the green channel for instance, you can begin to introduce more magenta into the tones by decreasing the curve. So dropping the green curve in the dark tones, and letting it settle back to normal towards the high tones will produce results like this:



Intentionally Removed the Image from here. It is not so important.

Suppressing the **green** channel in darks/mids adds a bit of **magenta** (click for original).

In isolation, these curves are fun to play with, but I think that perhaps walking through some actual examples of color toning/grading would help to illustrate what I’m talking about here. I’ll choose a couple of common toning examples to show what happens when you begin mixing all three channels up.

Color Toning/Grading

Orange and Teal Hell

I use the (*cinema film*) term *color grading* here because the first adjustment we will have a look at to illustrate curves is a horrible hollywood trend that is best described by [Todd Miro on his blog](#).

Grading is a term for color toning on film, and Todd’s post is a funny look at the prevalence of orange and teal in modern film palettes. So it’s worth a look just to see how silly this is (and hopefully to raise awareness of the obnoxiousness of this practice).

The general thought here is that caucasian skin tones trend towards orange, and if you have a look at a complementary color on the color wheel, you’ll notice that directly opposite orange is a teal color.



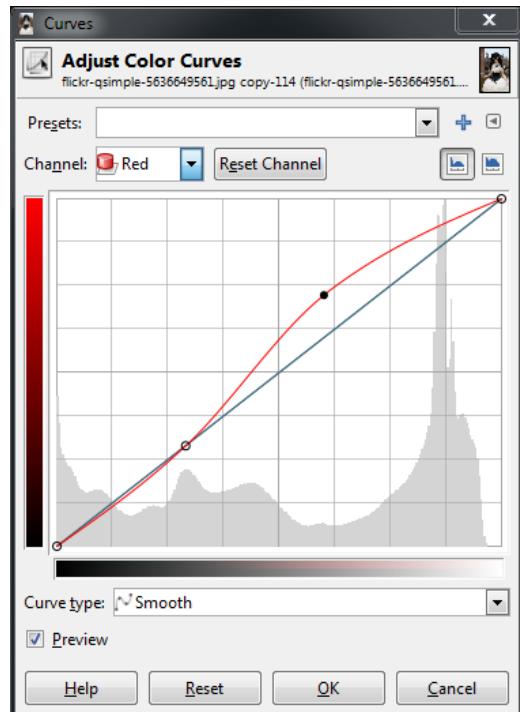
Screenshot from [Kuler](#) borrowed from Todd.

If you don’t already know about it, Adobe has online a fantastic tool for color visualization and palette creation called [Kuler](#) [Adobe Color CC](#). It lets you work on colors based on some classic rules, or even generate a color palette from images. Well worth a visit and a fantastic bookmark for fiddling with color.

So a quick look at the desired effect would be to keep/boost the skin tones into a sort of orange-y pinkish color, and to push the darker tones into a teal/cyan combination. (Colorists on films tend to use a Lift, Gamma, Gain model, but we’ll just try this out with our curves here).

Quick disclaimer - I am purposefully exaggerating these modifications to illustrate what they do. Like most things, moderation and restraint will go a long ways towards not causing your viewers eyeballs to bleed. *Remember - light touch!*

So I know that I want to see my skin tones head into an orange-ish color. In my image the skin tones are in the upper mids/low highs range of values, so I will start around there.



What I’ve done is put a point around the low midtones to anchor the curve closer to normal for those tones. This lets me fiddle with the red channel and to isolate it roughly to the mid and high tones only. The skin tones in this image in the red channel will fall toward the upper end of the mids, so I’ve boosted the reds there. Things may look a little weird at first:

Intentionally Removed the Image from here. It is not so important.

If you look back at the color wheel again, you’ll notice that between red and green, there is a yellow, and if you go a bit closer towards red the yellow turns

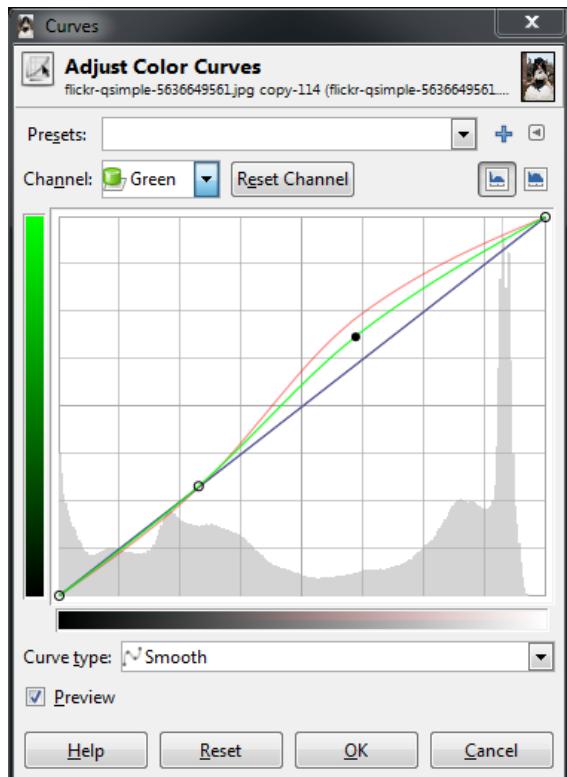
to more of an orange. What this means is that if we add some more green to those same tones, the overall colors will start to shift towards an orange.

So we can switch to the green channel now, put a point in the lower midtones again to hold things around normal, and slightly boost the green. Don't boost it all the way to the reds, but about 2/3rds or so to taste.

Conclusion

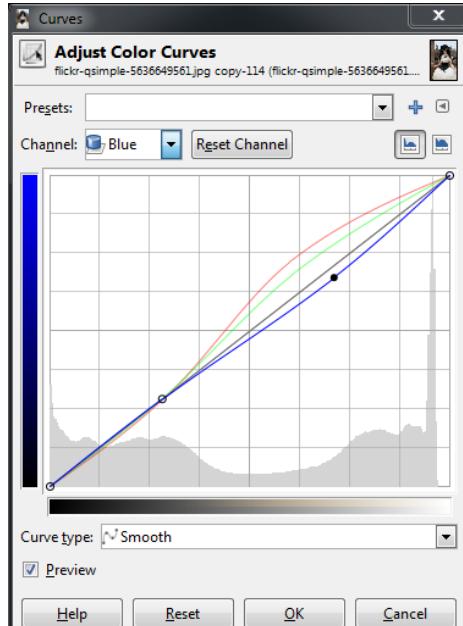
Remember; think about what the color curves represent in your image to help you achieve your final results. Begin looking at the different tonalities in your image and how you'd like them to appear as part of your final vision.

For even more fun - realize that the colors in your images can help to evoke emotional responses in the viewer, and adjust things accordingly. I'll leave it as an exercise for the reader to determine some of the associations between colors and different emotions.



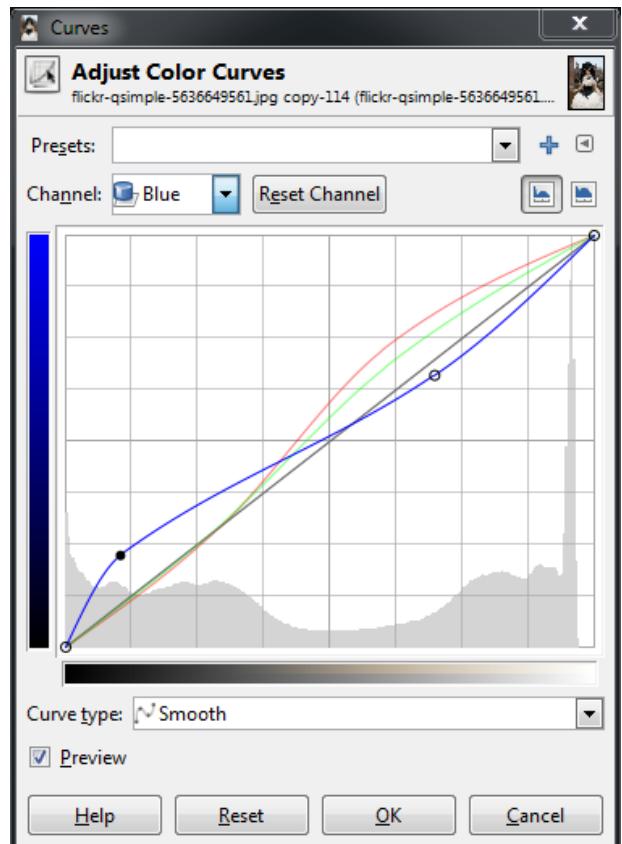
Intentionally Removed the Image from here. It is not so important.

This puts a little more red/orange-y color into the tones around the skin. You could further adjust this by perhaps including a bit more yellow as well. To do this, I would again put an anchor point in the low mid tones on the blue channel, then slightly drop the blue curve in the upper tones to introduce a bit of yellow.



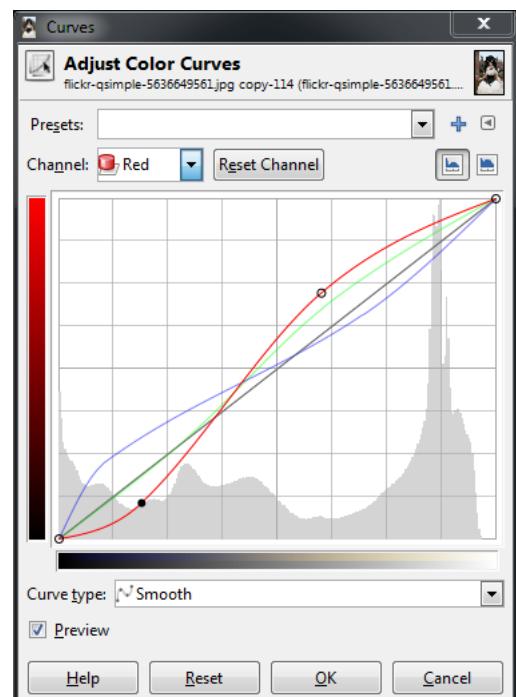
Intentionally Removed the Image from here. It is not so important.

Remember, we're experimenting here so feel free to try things out as we move along. I may consider the upper tones to be finished at the moment, and now I would want to look at introducing a more blue/teal color into the darker tones. I can start by boosting a bit of blues in the dark tones. I'm going to use the anchor point I already created, and just push things up a bit.



Intentionally Removed the Image from here. It is not so important.

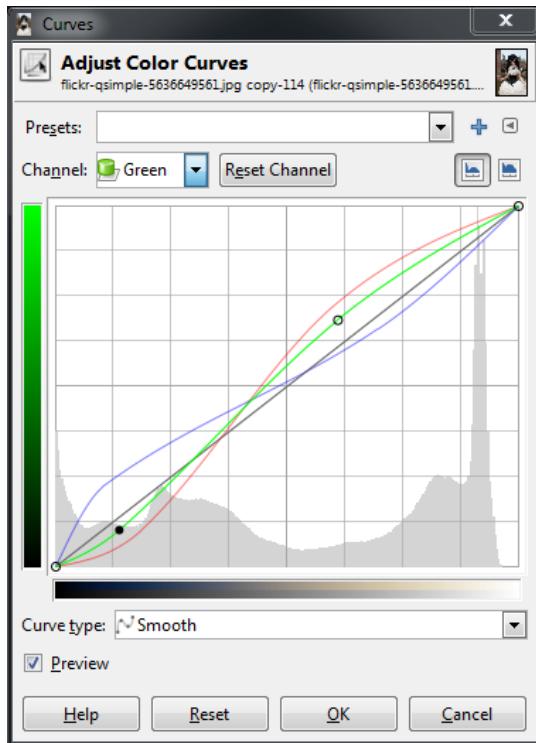
Now I want to make the darker tones a bit more teal in color. Remember the color wheel - **teal** is the absence of red - so we will drop down the red channel in the lower tones as well.



Finally to end up at something like this.

Intentionally Removed the Image from here. It is not so important.

And finally to push a very slight magenta into the dark tones as well, I'll push down the green channel a bit.



Intentionally Removed the Image from here. It is not so important.

After fooling around a bit - disgusting, isn't it? (click for original). I am exaggerating things here to illustrate a point. Please don't do this to your photos. :)

Conclusion

Remember, think about what the color curves represent in your image to help you achieve your final results. Begin looking at the different tonalities in your image and how you'd like them to appear as part of your final vision. For even more fun - realize that the colors in your images can help to evoke emotional responses in the viewer, and adjust things accordingly. I'll leave it as an exercise for the reader to determine some of the associations between colors and different emotions.

Intentionally Removed the Image from here. It is not so important.

If I wanted to go a step further, I could also put an anchor point up close to the highest values to keep the brightest parts of the image closer to a white instead of carrying over a color cast from our previous operations.

If your previous operations also darkened the image a bit, you could also now revisit the **Value** channel, and make modifications there as well. In my case I bumped the midtones of the image just a bit to brighten things up slightly.

